

CLAIMS

What is claimed is:

1. A method for maintaining cache coherency between a first controller and a peer controller, the method comprising:

updating cache data in a write-back cache memory in the first controller in response to read or write requests from an attached host;

buffering information regarding changes to the cache memory in the first controller that affect cache coherency of a mirrored cache memory in the peer controller; and

transmitting the buffered information and associated cache data from the first controller to the peer controller to maintain cache coherency between the cache memory of the first controller and that mirrored cache memory of the peer controller,

wherein the buffered information serves to maintain cache coherency without requiring a step of informing the peer controller of any cache flush operation performed in the first controller.

2. The method of claim 1

wherein the step of buffering further comprises:

filling a first buffer while a second buffer is transmitting to the peer controller; sensing completion of transmission of the second buffer to the peer controller;

and

filling the second buffer while the first buffer is transmitting to the peer controller.

3. The method of claim 2

wherein the step of transmitting further comprises:

transmitting the first buffer when any information is stored in the first buffer and the second buffer is not transmitting; and

transmitting the second buffer when any information is stored in the second buffer and the first buffer is not transmitting.

4. The method of claim 1

wherein the step of transmitting further comprises:

transmitting the buffered information and associated cache data using a single command structure.

5. The method of claim 4 wherein the step of transmitting further comprises: transmitting a SCSI Write Buffer command for transmitting the buffered information and the associated cache data.

6. The method of claim 1 wherein the step of buffering further comprises: generating a header information block indicative of a range of affected cache data corresponding to a read or write operation performed by the first controller; and adding the generated header information block to a filling buffer list of header information blocks.

7. The method of claim 6 wherein the step of transmitting further comprises: determining whether an active buffer list previously filled with header information blocks and associated cache data is presently transmitting to the peer controller; in response to determining that no active buffer list is presently transmitting, performing the steps of: designating the filling buffer list as the active buffer list; initiating transfer of the active buffer list to the peer controller; and designating another buffer list as the filling buffer list.

8. The method of claim 1 wherein the step of updating further comprises flushing cache data from the write-back cache memory in the first controller to thereby make room available in the write-back cache memory, and

wherein the step of buffering information further comprises buffering information to invalidate cache data in the mirrored cache memory in the peer controller that corresponds to the dirty data flushed by the first controller if the step of updating is responsive to a read request from an attached host.

9. A system comprising:

- a first storage controller having a cache memory;
- a second storage controller having a mirrored cache memory; and
- a communication channel coupling the first controller and the peer controller for exchanging cache coherency information,

wherein the first controller further comprises:

- a filling buffer list for accumulating header information regarding changes in the cache memory of the first controller;
- an active buffer list for transmitting previously accumulated header information regarding changes in the cache memory of the first controller from the first controller to the peer controller; and
- a first coherency element coupled to the filling buffer list and coupled to the active buffer list to generate header information regarding changes in the cache memory of the first controller resulting from performing read and write operations and to store the generated header information in the filling buffer and to transmit the active buffer to the peer controller,

wherein the second controller further comprises:

- a second coherency element for receiving the transmitted header information from the first coherency element to update the mirrored cache memory accordingly, and

wherein the first controller forwards no cache data to the second controller in response to periodic cache flush operations performed by the first controller.

10. A method for maintaining cache coherency between a cache memory in a first controller and a mirror cache memory in a peer controller, the method comprising the steps of:

- generating in the first controller update meta-data regarding dirty data in the cache memory of the first controller altered in response to write requests received from an attached host system;
- generating in the first controller invalidation meta-data regarding dirty data in the cache memory of the first controller flushed by the first controller in response to read or write requests received from an attached host system;
- periodically transmitting the generated meta-data and any corresponding dirty data from the first controller to the peer controller to permit the peer controller to

update its mirrored cache memory in accordance with the received meta-data and any corresponding cache data.

11. The method of claim 10 wherein the steps of generating each include generating meta-data into a currently filling buffer as a previously filled buffer is being transmitted and wherein the step of periodically transmitting includes the step of swapping the currently filling buffer and previously filled buffer after completion of the transmission of the previously filled buffer.

12. A system for maintaining cache coherency between a first controller and a peer controller, the system comprising:

updating means for updating cache data in a write-back cache memory in the first controller in response to read or write requests from an attached host;

buffering means for buffering information regarding changes to the cache memory in the first controller that affect cache coherency of a mirrored cache memory in the peer controller; and

transmitting means for transmitting the buffered information and associated cache data from the first controller to the peer controller to maintain cache coherency between the cache memory of the first controller and that mirrored cache memory of the peer controller,

wherein the buffered information serves to maintain cache coherency without requiring a step of informing the peer controller of any cache flush operation performed in the first controller.

13. The system of claim 12

wherein the buffering means further comprises:

means for filling a first buffer while a second buffer is transmitting to the peer controller;

means for sensing completion of transmission of the second buffer to the peer controller; and

means for filling the second buffer while the first buffer is transmitting to the peer controller.

14. The system of claim 13
wherein the transmitting means further comprises:
means for transmitting the first buffer when any information is stored in the first buffer and the second buffer is not transmitting; and
means for transmitting the second buffer when any information is stored in the second buffer and the first buffer is not transmitting.

15. The system of claim 12
wherein the transmitting means further comprises:
means for transmitting the buffered information and associated cache data using a single command structure.

16. The system of claim 15 wherein the means for transmitting further comprises:
means for transmitting a SCSI Write Buffer command for transmitting the buffered information and the associated cache data.

17. The system of claim 12 wherein the buffering means further comprises:
means for generating a header information block indicative of a range of affected cache data corresponding to a read or write operation performed by the first controller; and
means for adding the generated header information block to a filling buffer list of header information blocks.

18. The system of claim 17
wherein the transmitting means further comprises:
means for determining whether an active buffer list previously filled with header information blocks and associated cache data is presently transmitting to the peer controller;
means for determining that no active buffer list is presently transmitting;
means responsive to the means for determining for designating the filling buffer list as the active buffer list;
means responsive to the means for determining for initiating transfer of the active buffer list to the peer controller; and

means responsive to the means for determining for designating another buffer list as the filling buffer list.

19. The system of claim 12

wherein the updating means further comprises means for flushing cache data from the write-back cache memory in the first controller to thereby make room available in the write-back cache memory, and

wherein the buffering means further comprises means for buffering information to invalidate cache data in the mirrored cache memory in the peer controller that corresponds to the dirty data flushed by the first controller if the updating means is operating in response to a read request from an attached host.